arXiv:1609.05194v1 [math.CO] 16 Sep 2016

The Bradley–Terry condition is L_1–testable

Agelos Georgakopoulos*[1] and Konstantinos Tyros†[2]

[1]Mathematics Institute, University of Warwick, CV4 7AL, UK,
[2]Department of Mathematics, Koç University, Saryer, Istanbul, 34450, Turkey

Abstract

We provide an algorithm with constant running time that given a weighted tournament T, distinguishes with high probability of success between the cases that T can be represented by a Bradley–Terry model, or cannot even be approximated by one. The same algorithm tests whether the corresponding Markov chain is reversible.

Keywords: Bradley–Terry model, property testing, L_1-tester, reversible Markov chain, stochastic tournament

1 Introduction

Suppose we have a set S of individuals participating in a tournament, and for each pair x, y in S we have assigned a probability p_{xy} that x beats y (so that $p_{yx} = 1 - p_{xy}$). The Bradley–Terry model [6] seeks to assign real numbers $a(x) > 0$ to these individuals so that

$$\frac{p_{xy}}{p_{yx}} = \frac{a(x)}{a(y)}. \tag{1}$$

holds for every x, y in S; equivalently, we have $p_{xy} = \frac{a(x)}{a(x)+a(y)}$.

This statistical model, which was already considered by Zermelo [15], is used in many practical applications where individuals can be compared in pairs, and the probabilities p_{xy} are often measured empirically[1].

Although it is convenient to represent the $\binom{|S|}{2}$ parameters p_{xy} with the just $|S|$ parameters $a(x)$, it is easy to see that not all tournaments can be approximated well in this way: consider for example a tournament containing three vertices x_0, x_1, x_2 such that x_i beats $x_{i+1 \pmod 3}$ with probability 90%. The model has a widespread use nevertheless, and the algorithm used for computing the $a(x)$ is known to converge even for such tournaments [8].

*Supported by EPSRC grant EP/L002787/1, and by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No 639046).

†Supported by ERC grant 306493.

[1]en.wikipedia.org/wiki/Bradley-Terry_model

Thus it is important to know when a Bradley–Terry model approximates the probabilities p_{xy} well-enough. An obvious but rather inefficient approach is to first compute the $a(x)$, and then test them against (1) for every pair x, y.

The aim of this paper is to provide a fast randomized algorithm that given a tournament T, distinguishes with high probability of success between the cases that T can be represented by a Bradley–Terry model, or cannot even be approximated by one. The running time of our algorithm is independent of the size of T; it only depends on the approximation tolerance and the desired probability of success. It is assumed here that —unlike in statistical setups— our algorith can query the exact value p_{xy} in one step.

More consisely, we prove that the Bradley–Terry condition is L_1–testable in the sense of [5].

2 Preliminaries

2.1 Property testing

The notion of property testing was established by Goldreich, Goldwasser and Ron [10] for graph properties, following a similar notion of Rubinfeld and Sudan [14] for program testing in computer science. Since then the concept has received a lot of attention in various contexts; see [12]. We recall the following standard definitions from [2].

A *property* $\mathcal{P}$ is a family of (undirected) graphs closed under isomorphism. A graph G with n vertices is *ϵ-far* from satisfying $\mathcal{P}$ if one must add or delete at least ϵn^2 edges in order to turn G into a graph in $\mathcal{P}$.

An *ϵ-tester* for $\mathcal{P}$ is a randomized algorithm, which given the ability to check whether there is an edge between a given pair of vertices, distinguishes with probability at least $2/3$ between the cases G satisfies $\mathcal{P}$ and G is ϵ-far from satisfying $\mathcal{P}$. Such an ϵ-tester is *one-sided* if, whenever G satisfies $\mathcal{P}$, the ϵ-tester determines this with probability 1. A property $\mathcal{P}$ is *strongly-testable* if for every fixed $\epsilon > 0$ there exists a one-sided ϵ-tester for $\mathcal{P}$ whose query complexity is bounded only by a function of ϵ, which is independent of the size of the input graph. Call a property $\mathcal{P}$ *easily testable* if it is strongly testable with a one-sided ϵ-tester whose query complexity is polynomial in ϵ^{-1}, and otherwise call $\mathcal{P}$ *hard*.

It is easy to adapt the above definitions to the class of edge-weighted graphs: say that a weighted graph (G, w) with n vertices is *ϵ-far* from satisfying $\mathcal{P}$ if the total change in w required in order to turn (G, w) into a weighted graph in $\mathcal{P}$ is at least ϵn^2. All other definitions can be applied verbatim and, following [5], we say that a property is (easily/strongly) L_1–testable if it complies with the above definitions when we use this new notion of ϵ-far.

2.2 Stochastic tournaments

A *stochastic tournament* is a pair (K, w) where K is a *tournament*, i.e. a complete graph with n vertices in which every edge is given a direction, and $w : E(K) \to [0, 1]$ is an assignment of weights to the (directed) edges of K. Intuitively, we think of $w(xy)$ as the probability that x beats y in a game between them. The *matrix* $P = P(K, w) = (p_{xy})$ of the stochastic tournament (K, w) is the square matrix with rows and columns indexed by the vertices of K such that

$p_{xy} = w(xy) = 1 - p_{yx}$ whenever the edge xy, directed from x to y, belongs to $E(K)$. A stochastic tournament (K, w) satisfies the *Bradley–Terry condition*, if there exists a family of positive real numbers (a_x) indexed by the vertex set of K such that condition (1) holds. Moreover, a stochastic tournament corresponds to a Markov chain by considering the matrix $Q = Q(K, w) = (q_{xy})_{x,y \in V(K)}$ defined by

$$q_{xy} = \begin{cases} p_{xy}/n, & \text{if } x \neq y \\ 1 - \frac{1}{n} \sum_{z \neq x} p_{xz}, & \text{if } x = y \end{cases} \tag{2}$$

Then Q defines indeed a Markov chain by interpreting q_{xy} as the transition probability from vertex x to vertex y. Let us recall that such a Markov chain with transition matrix $(q_{xy})_{x,y \in V(K)}$ is called reversible, if there exists a probability distribution (π_x) on the vertex set of K such that for every pair of nodes x, y we have

$$\pi_x q_{xy} = \pi_y q_{yx}. \tag{3}$$

All these notions are related to each other as it is described by the following partly folklore proposition.

Proposition 2.1. *Let (K, w) be a stochastic tournament. Then the following conditions are equivalent*

(i) (K, w) corresponds to a Bradley–Terry model;

(ii) The Markov chain corresponding to Q is reversible;

(iii) Every triangle of (K, w) is balanced, i.e. for every triple x, y, z in $V(K)$, we have $p_{xy}p_{yz}p_{zx} = p_{xz}p_{zy}p_{yx}$;

(iv) There is a vertex r in $V(K)$ such that every triangle containing r is balanced.

The equivalence of (i) and (ii) was noticed by Rajkumar & Agarwal [13, Section 4], while the equivalence of (ii) and (iii) is essentially Kolmogorov's criterion.

Proof. (i) = (ii): Let the stationary distribution $\pi(x)$ of a vertex equal the inverse $1/a(x)$ of its Bradley–Terry score. Then equivalence follows easily from the definition of reversibility: we have $\pi(x)Q(xy) = \pi(y)Q(yx)$ if and only if $P(xy)/P(yx) = a(x)/a(y)$.

(iv) $\to$ (i): Given such a vertex r, let $a(r) = 1$. For every other vertex y, let $a(y) := p_{yr}/p_{ry}$. We claim that these numbers satisfy the Bradley–Terry condition, i.e. $p_{zy}/p_{yz} = a(z)/a(y)$ for every y, z in $V(K)$. Indeed, note that $\frac{p_{ry}}{p_{yr}} \frac{p_{zr}}{p_{rz}} = \frac{1}{a(y)} \frac{a(z)}{1}$ by the definition of a. Combining this with (iv) applied to the triangle r, y, z results in cancellations yielding the desired $p_{zy}/p_{yz} = a(z)/a(y)$.

(i) $\to$ (iii): We have $\frac{p_{xy}p_{yz}p_{zx}}{p_{xz}p_{zy}p_{yx}} = \frac{a(x)a(y)a(z)}{a(x)a(z)a(y)} = 1$ by applying (1) three times. Finally, the implication (iii) $\to$ (iv) is trivial completing the proof. □

Theorem 2.2. *Each of the four conditions of Proposition 2.1 is easily L_1–testable.*

As observerd in [5, Fact 1.1.], L_1–testability implies L_p testability for all $p > 1$. The proof of Theorem 2.2 is based on the balanced triangle condition of Proposition 2.1 (iii), and parallels a classical proof of testability of triangle-freeness for dense graphs. For this we use a continuous analogue of the triangle removal lemma, which we prove in the next section (Corollary 3.3).

In Section 4 we show that the notion of being ϵ-far from satisfying these conditions corresponds to natural approximate versions of the Bradley–Terry condition and reversibility.

3 Proof of Theorem 2.2

Let us observe that given any stochastic tournament (K, w) and any triangle T with nodes x, y and z, we can make T balanced by changing the value w on any of its edges. More precisely, T can be made balanced by changing, for instance, the value of w only on the edge xy yielding p_{xy} to be equal to $\frac{p_{zy}p_{xz}}{p_{zy}p_{xz}+p_{yz}p_{zx}}$.

Definition 3.1. *Let (K, w) be a stochastic tournament and T a triangle with nodes x, y and z. We define the discrepancy of T, denoted by $\mathrm{disc}_{(K,w)}(T)$ to be the maximum of the changes of w required on a single edge of T to turn it into a balanced triangle, that is,*

$$\mathrm{disc}_{(K,w)}(T) = \max\left\{|\alpha|, |\beta|, |\gamma|\right\} \tag{4}$$

where $\alpha = p_{xy} - \frac{p_{zy}p_{xz}}{p_{zy}p_{xz}+p_{yz}p_{zx}}$, $\beta = p_{yz} - \frac{p_{yx}p_{xz}}{p_{yx}p_{xz}+p_{xy}p_{zx}}$ and $\gamma = p_{zx} - \frac{p_{yx}p_{zy}}{p_{yx}p_{zy}+p_{xy}p_{yz}}$.

We will omit the subscript from the symbol $\mathrm{disc}_{(K,w)}$ whenever the stochastic tournament (K, w) is clear from the context. Let us observe that T is balanced if and only if $\mathrm{disc}(T) = 0$.

The following lemma is based on the observation that we can make any stochastic tournament reversible by modifying the edges that do not belong to any fixed spanning tree; see Lemma 5.1. Here we work with the spanning tree consisting of all edges incident with a fixed vertex r.

Lemma 3.2. *Let (K, w) be a stochastic tournament and r a node in K. Also let $\mathcal{C}$ be the collection of all triangles containing r and for each T in $\mathcal{C}$ let e_T be the edge of T that does not contain r. Then there is a weighting w' satisfying the following:*

(i) the stochastic tournament (K, w') is reversible,

(ii) $w'(e) = w(e)$ for all e in $E(K)$ that do not belong to $\{e_T : T \in \mathcal{C}\}$ and

(iii) $|w'(e_T) - w(e_T)| \leqslant \mathrm{disc}_{(K,w)}(T)$ for all T in $\mathcal{C}$.

Proof. By the definition of $\mathrm{disc}_{(K,w)}(T)$, for each T in $\mathcal{C}$, we can assign to e_T a new weight $w'(e_T)$ making T balanced and satisfying $|w(e_T) - w'(e_T)| < \mathrm{disc}_{(K,w)}(T)$. Since for every T in $\mathcal{C}$ we know that T is the only triangle in $\mathcal{C}$ containing e_T, we deduce that every triangle in $\mathcal{C}$ is balanced. The result follows by Proposition 2.1. □

Averaging this over all vertices r of K, we deduce

Corollary 3.3. *Let (K, w) be a stochastic tournament satisfying*

$$\sum_T \mathrm{disc}(T) \leqslant \epsilon \binom{|K|}{3},$$

for some $\epsilon > 0$, where the sum is taken over all triangles. Then there is a weighing w' such that the tournament (K, w') is reversible, and

$$\sum_{e \in E(K)} |w(e) - w'(e)| < \epsilon \binom{|K|}{2}. \tag{5}$$

Proof. Let $\mathcal{T}$ be the collection of all triangles in K and for every node r in K, let $\mathcal{C}_r$ be the collection of all triangles containing r. Let $n = |K|$. Observe that

$$\sum_{r \in V(K)} \sum_{T \in \mathcal{C}_r} \mathrm{disc}(T) = 3 \sum_{T \in \mathcal{T}} \mathrm{disc}(T) < 3\epsilon \binom{n}{3} < n\epsilon \binom{n}{2} \tag{6}$$

and therefore there is some r in $V(K)$ such that $\sum_{T \in \mathcal{C}_r} \mathrm{disc}(T) < \epsilon \binom{n}{2}$. The result follows from Lemma 3.2. □

We now use the last result to prove our main theorem.

Proof of Theorem 2.2. We consider the following randomised algorithm. Sample $f(\epsilon) = \lceil -\log_{1-\epsilon} 3 \rceil$ triangles of K independently and uniformly at random, and check whether they are balanced. Answer 'yes' if all these triangles are balanced, and answer 'no' otherwise.

Clearly, if (K, w) satisfies the conditions of Proposition 2.1, our algorithm responds 'yes' with probability one by item (iii). On the other hand, assuming that (K, w) is ϵ-far from satisfying these conditions, that is, (5) fails for every w' such that (K, w') satisfies the Bradley–Terry condition, Corollary 3.3 implies

$$\sum_T \mathrm{disc}(T) > \epsilon \binom{|K|}{3}.$$

Since $0 \leqslant \mathrm{disc}(T) \leqslant 1$, letting B be the set of the unbalanced triangles, we obtain $|B| \geqslant \epsilon \binom{|K|}{3}$. By the choice of $f(\epsilon)$, our algorithm responds 'no' with probability at least $2/3$; indeed, each of our sampled triangles is in B with probability at least ϵ. Thus the probability that none of them is in B is at most $(1-\epsilon)^{f(\epsilon)} \leqslant 1/3$. It is not hard to check that $f(\epsilon) < c\epsilon^{-1}$ for some constant c (in fact we can take $c = 2$), and so our property is easily L_1–testable. □

Remark 1: The above proof parallels the classical proof of testability of triangle-freeness for dense graphs, with Corollary 3.3 playing the role of the triangle removal lemma. But as the same ϵ appears in the condition and the conclusion of Corollary 3.3, the conditions of Proposition 2.1 are easily testable contrary to triangle-freeness which is known to be hard [7, p. 4].

Remark 2: Our proof in fact shows the slightly stronger fact that the conditions of Proposition 2.1 are easily L_0–testable in the sense of [5, Fact 1.1.], but we chose to work with L_1 as it is more natural in our setup. Indeed, we could have defined $\mathrm{disc}(T)$ to be 1 if T is unbalanced and 0 otherwise, and the rest of our proof could be used verbatim.

4 Approximate versions of Proposition 2.1

Let (K, w) be a stochastic tournament and ϵ a real with $0 < \epsilon \leqslant 1$. We will say that (K, w) is an ϵ-approximate Bradley–Terry model if there is a map $a : V(K) \to \mathbb{R}_{>0}$ such that

$$(1+\epsilon)^{-1}\frac{a(x)}{a(x)+a(y)} \leqslant p_{xy} \leqslant (1+\epsilon)\frac{a(x)}{a(x)+a(y)}$$

for all pairs of distinct vertices x, y in $V(K)$. We will call the map a an ϵ-approximate Bradley–Terry score. Moreover, a triangle with nodes x, y and z will be called ϵ-balanced if

$$(1+\epsilon)^{-1} \leqslant \frac{p_{xy}p_{yz}p_{zx}}{p_{yx}p_{zy}p_{xz}} \leqslant 1+\epsilon.$$

Proposition 4.1. *Let (K, w) be a stochastic tournament and ϵ a real with $0 < \epsilon \leqslant 1$. Also let r be a node in K and assume that every triangle containing r is ϵ-balanced. Then (K, w) is an ϵ-approximate Brandley–Terry model.*

Proof. We define $a : V(K) \to \mathbb{R}_{>0}$ setting $a(r) = 1$ and for every node y in K with $y \neq r$ we set $a(y) = \frac{p_{yr}}{p_{ry}}$. We claim that a is an ϵ-approximate Bradley–Terry score. Indeed, it follows directly from the definition of a that for every node y in K with $y \neq r$, we have

$$p_{ry} = \frac{a(r)}{a(r)+a(y)} \text{ and } p_{yr} = \frac{a(y)}{a(r)+a(y)}.$$

Let x, y be two distinct nodes in K different from r. Since the triangle consisting of the nodes x, y and r is ϵ-balanced, we obtain

$$(1+\epsilon)^{-1} \leqslant \frac{p_{yx}}{p_{xy}} \cdot \frac{p_{ry}}{p_{yr}} \cdot \frac{p_{xr}}{p_{rx}} = \frac{p_{yx}}{p_{xy}} \cdot a(x) \cdot \frac{1}{a(y)} \leqslant 1+\epsilon \tag{7}$$

and therefore

$$(1+\epsilon)^{-1}\frac{a(y)}{a(x)} \leqslant \frac{p_{yx}}{p_{xy}} \leqslant (1+\epsilon)\frac{a(y)}{a(x)}. \tag{8}$$

By (8), we get that

$$\begin{aligned}(1+\epsilon)^{-1}\Big(1+\frac{a(y)}{a(x)}\Big) < 1 + (1+\epsilon)^{-1}\frac{a(y)}{a(x)} &\leqslant 1 + \frac{p_{yx}}{p_{xy}} \\ < 1 + (1+\epsilon)\frac{a(y)}{a(x)} &\leqslant (1+\epsilon)\Big(1+\frac{a(y)}{a(x)}\Big)\end{aligned} \tag{9}$$

and therefore

$$(1+\epsilon)^{-1}\frac{1}{1+\frac{a(y)}{a(x)}} < \frac{1}{1+\frac{p_{yx}}{p_{xy}}} < (1+\epsilon)\frac{1}{1+\frac{a(y)}{a(x)}}. \tag{10}$$

Finally, since

$$\frac{a(x)}{a(x)+a(y)} = \frac{1}{1+\frac{a(y)}{a(x)}} \text{ and } p_{xy} = \frac{1}{1+\frac{p_{yx}}{p_{xy}}},$$

by inequality (10), the result follows. □

Let (K, w) be a stochastic tournament and ϵ a real with $0 < \epsilon \leqslant 1$. We will say that (K, w) is ϵ-approximately reversible, if there is a probability distribution $(\pi_x)_x$ such that for every pair of distinct nodes x and y in K we have

$$(1+\epsilon)^{-1} \leqslant \frac{\pi_x p_{xy}}{\pi_y p_{yx}} \leqslant 1+\epsilon.$$

Proposition 4.2. *Let (K, w) be a stochastic tournament and ϵ a real with $0 < \epsilon \leqslant 1$. Assume that (K, w) is an ϵ-approximate Brandley–Terry model. Then (K, w) is 3ϵ-approximately reversible.*

Proof. Let a be an ϵ-approximate Brandley–Terry score. Without loss of generality we may assume that $\sum_x 1/a(x) = 1$, where the sum is taken over all nodes of K. For every node x of K we set $\pi_x = 1/a(x)$. We claim that $(\pi_x)_x$ witnesses that (K, w) is ϵ-approximately reversible. Indeed, let x and y be to distinct nodes of K. First observe that since a is an ϵ-approximate Brandley–Terry score, we have that

$$(1+\epsilon)^{-1}\frac{a(x)}{a(x)+a(y)} \leqslant p_{xy} \leqslant (1+\epsilon)\frac{a(x)}{a(x)+a(y)} \tag{11}$$

and

$$(1+\epsilon)^{-1}\frac{a(y)}{a(x)+a(y)} \leqslant p_{yx} \leqslant (1+\epsilon)\frac{a(y)}{a(x)+a(y)}. \tag{12}$$

Thus, we have that

$$\frac{\pi_x p_{xy}}{\pi_y p_{yx}} = \frac{a(x)^{-1} p_{xy}}{a(y)^{-1} p_{yx}} \leqslant \frac{1+\epsilon}{(1+\epsilon)^{-1}} \leqslant 1+3\epsilon. \tag{13}$$

Similarly it follows that $\frac{\pi_x p_{xy}}{\pi_y p_{yx}} \geqslant (1+3\epsilon)^{-1}$. □

Finally, it is immediate that if a stochastic tournament (K, w) is a ϵ-approximately reversible then every triangle is 7ϵ-balanced.

5 Balanced cycles

In this section we show that if all cycles in any basis of the cycle space of our stochastic tournament (K, w) are balanced, then every cycle of K is balanced.

Lemma 5.1. *Let K be a tournament and S an (undirected) spanning tree of K. Then every weighting $\hat{w} : E(S) \to \mathbb{R}_{>0}$ can be extended to a weighting $w : E(K) \to \mathbb{R}_{>0}$ of K such that the Markov chain defined by $Q(K, w)$ is reversible.*

As we observe below, Lemma 5.1 can be thought of as a special case of a more general result Lemma 5.2.

Proof of Lemma 5.1. We will define the extension w of $\hat{w}$ by first defining the stationary measure π of the Markov chain of $Q(K, w)$.

For this, fix a vertex r of K, and let $\pi(r) = 1$ (any positive value would do). For each neighbour y of r in S, we let $\pi(y) = \pi(r)\hat{w}(ry)/(1 - \hat{w}(ry))$, where we set $\hat{w}(ry) = 1 - \hat{w}(yr)$ if the ry–edge is directed from y to r.

We proceed recursively to assign a value $\pi(x)$ to each neighbour x of y except r, by the same formula: $\pi(x) = \pi(y)\hat{w}(yx)/(1-\hat{w}(yx))$. This defines π. Now for every chord $e = xy$ of S, we let $w(xy)$ be the unique solution to $\frac{w(xy)}{1-w(xy)} = \frac{\pi(y)}{\pi(x)}$, that is, $w(xy) = c/(1+c)$ where $c = \frac{\pi(y)}{\pi(x)}$. It follows that the measure π is stationary for $Q(K,w)$, since we have $\pi(x)q_{xy} = \pi(x)\frac{w(xy)}{n} = \pi(y)\frac{w(yx)}{n} = \pi(y)q_{yx}$ by the definitions. □

Call a cycle C of K *balanced*, if $\lambda(\overrightarrow{C}) := \Pi_{xy\in\overrightarrow{C}}\frac{p_{xy}}{p_{yx}} = 1$, where $\overrightarrow{C}$ denotes any of the two possible cyclic orientations of C.

Lemma 5.2. *Let (K,w) be a stochastic tournament. Let B be a basis of the cycle space $\mathcal{C}_\mathbb{Z}$ of K. If every element of B is balanced, then every element of $\mathcal{C}$ is balanced (and hence the Markov chain defined by $Q(K,w)$ is reversible).*

Proof. It is straightforward to check that if $\overrightarrow{C}, \overrightarrow{D} \in \mathcal{C}_\mathbb{Z}$ then $\lambda(\overrightarrow{C}+\overrightarrow{D}) = \lambda(\overrightarrow{C})\lambda(\overrightarrow{D})$ by the definition of λ. Thus any element of $\mathcal{C}_\mathbb{Z}$ generated by a balanced set is balanced, and the result follows. By Proposition 2.1, $Q(K,w)$ is reversible in this case. □

We could have deduced Lemma 5.1 from the last result as follows. For every chord $e = xy$ of T, we can assign a value $w(xy)$ such that the fundamental cycle of e with respect to T becomes balanced. The result then follows from the well-known fact that the fundamental cycles of any spanning tree generate the cycle space [9].

6 Open Problems

Suppose that instead of a tournament K we have an arbitrary directed graph G, and an assignment of weights $w : E(G) \to (0,1)$ to the edges of G. Then the Markov chain corresponding to Q is still well-defined if we set $p_{xy} = 0$ whenever x, y do not form an edge of G. Thus we can ask if the Markov chain is reversible. Likewise, we can generalise the Bradley–Terry condition by demanding (1) only when x, y form an edge. We expect that using a version of the Szemeredi regularity lemma it is possible to show that these properties are L_1–testable for arbitrary G following the approach of [1] or [4], which characterises the (unweighted) graph properties that are testable. However, this would give far worse bounds than those we obtained for tournaments, which naturally leads to the following question.

Problem 6.1. *Are reversibility and the Bradley–Terry condition easily L_1–testable for arbitrary directed graphs?*

We remark that the problem of characterising easily testable graph properties is open [3].

The notion of testability has been adapted to sparse graphs [11]; the definition is the same, except that we say that G is *ϵ-far* from having a property $\mathcal{P}$ if one must add or delete at least ϵn (rather than ϵn^2) edges in order to turn G into a graph having $\mathcal{P}$.

Problem 6.2. *Are reversibility and the Bradley–Terry condition L_1–testable in the sparse sense?*

Next, we propose a generalisation of the Bradley–Terry condition. Although it also makes sense for arbitrary graphs, we will formulate it for tournaments for simplicity.

Let us say that a stochastic tournament (K, w) has *Bradley–Terry dimension* (or *BT–dimension* for short) d, if there is a family $(a_i)_{1\leq i\leq d}$ of d functions $a_i : V(K) \to \mathbb{R}_{>0}$ such that, for every $xy \in V(K)$, we have

$$p_{xy} = \frac{1}{d} \sum_{1\leq i\leq d} \frac{a_i(x)}{a_i(x) + a_i(y)}, \tag{14}$$

and this is not true for any family of $d-1$ functions. Here p_{xy} is determined by w as explained in Section 2.2. Thus (K, w) has BT–dimension 1 if and only if it satisfies the Bradley–Terry condition.

Intuitively, if (K, w) has BT–dimension d, then we can represent the probabilities p_{xy} via a set of d games, such that each player $x \in V(K)$ has a strength $a_i(x)$ in game i, and when players x, y compete, one of these d games is chosen uniformly at random, and x, y play a match of that game.

Problem 6.3. *Is the property of having BT–dimension d L_1–testable for $d > 1$?*

One can think of BT–dimension as a generalisation of reversibility for Markov chains due to Proposition 2.1. It would be interesting to extend properties known for reversible Markov chains to Markov chains with bounded BT–dimension.

Acknowledgement

We would like to thank Oded Goldreich for several suggestions.

References

[1] N. Alon, E. Fischer, E. Newman, and A. Shapira. A Combinatorial Characterization of the Testable Graph Properties: It's All About Regularity. In *SIAM Journal on Computing (Special Issue of STOC'06)*, volume 39, pages 143–167, 2009.

[2] N. Alon and J. Fox. Easily testable graph properties. *Comb., Probab. Comput.*, 24:646–657, 2015.

[3] N. Alon and J. Fox. Easily Testable Graph Properties. *Comb., Probab. Comput.*, 24:646–657, 2015. Special Issue 04.

[4] N. Alon and A. Shapira. Every Monotone Graph Property Is Testable. *SIAM Journal on Computing*, 38(2):505–522, 2008.

[5] P. Berman, S. Raskhodnikova, and G. Yaroslavtsev. Lp-testing. In *46th STOC*, pages 164–173, 2014.

[6] Ralph Allan Bradley and Milton E. Terry. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39(3/4):324–345, 1952.

[7] David Conlon and Jacob Fox. Graph removal lemmas. In *Surveys in Combinatorics 2013*, LMS Lecture Note Series. Cambridge University Press, 2013.

[8] H.A. David. *The Method of Paired Comparisons*. Hodder Arnold, 1988.

[9] Reinhard Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005. Electronic edition available at:
`http://www.math.uni-hamburg.de/home/diestel/books/graph.theory`.

[10] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its applications to learning and approximation. *J. ACM*, 45:653–750, 1998.

[11] O. Goldreich and D. Ron. Property Testing in Bounded Degree Graphs. *Algorithmica*, 32:302–343, 2002.

[12] Oded Goldreich. Introduction to Property Testing. Monograph in preparation. http://www.wisdom.weizmann.ac.il/ oded/pt-intro.html.

[13] A. Rajkumar and S. Agarwal. A Statistical Convergence Perspective of Algorithms for Rank Aggregation from Pairwise Data. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 118–126, 2014.

[14] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM J. Comput.*, 25:252–271, 1996.

[15] E. Zermelo. Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 29(1):436–460, 1929.